ICL UT INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY of TENNESSEE

2012/2013 REPORT

INNOVATIVE COMPUTING LABORATORY **2012/2013** REPORT

EDITED BY **Sam Crawford** DESIGNED BY **David Rogers**

INNOVATIVE COMPUTING LABORATORY

THE UNIVERSITY of TENNESSEE

2012/2013 REPORT

Jack Dongarra
Director, ICL

FROM TH

Like other successful academic research groups, the Innovative Computing Laboratory participates in and reaches out to scientific and engineering communities that are global in scope; and yet we all recognize that the health and vitality of our local environment—our university, our state, our region—plays an essential role in our achievements. The reason is simple: the quality of the people that we attract and retain is the single most important factor in the success of competitive research organizations such as ours; and the type of creative individuals we depend on, in turn, see the support of a rich and vibrant local environment as critical to their ability to make meaningful contributions to their fields and to achieve important personal and professional goals. So as we end 2012 and begin a new year, preparing with the rest of our global community to make the multi-year climb to the new era of Exascale computing, it seems appropriate in this report to highlight the people of ICL, their achievements over the past year, and the way they have worked cooperatively to help put and keep the University of Tennessee and the East Tennessee region at the forefront of high performance scientific computing.

The chart of ICL alumni and major ICL projects that introduces our report is a reminder that we are building on more than two decades of remarkable creativity and dedication of the people who have come to ICL and UTK. The names of the projects that they have helped lead—PVM, MPI, LAPACK, ScaLAPACK, BLAS, ATLAS, Netlib, NHSE, TOP500, PAPI, NetSolve, Open MPI, FT-MPI, the HPC Challenge, and the LINPACK Benchmark—are familiar to HPC users around the world. And the software that they have produced now

E DIRECTOR

provides critical research infrastructure for hundreds of thousands, if not millions of users. But if you look at the probable path to Exascale computing, and survey the set of revolutionary problems it presents (e.g., the need to exploit billion way parallelism, order of magnitude increases in the number of execution faults, unprecedented constraints on energy use, and so on), it becomes clear that the most challenging years for ICL research may still lie in front of us.

We are fortunate, therefore, that today's ICL brings together one of the most talented and most experienced teams we have ever had. In the areas of numerical libraries, distributed computing, and performance monitoring and benchmarking, they have created and are leading ground breaking projects—PLASMA, MAGMA, PaRSEC, PAPI-V, Open MPI, to name a few—that are directly targeting some of the central problems in HPC today. Moreover, during the last several years, we have improved our ability to support these efforts. With generous support from the University, we have been able to build up our Center for Information Technology Research (CITR) into a first class research support organization, dedicated to maximizing the ability of our research teams to make new discoveries and innovations, develop original projects and execute them well, and improve the resulting flow of publications, software, and community engagement.

In terms of our local community, the coming year will offer us the opportunity to make important contributions on two key fronts. First, UTK and the National Institute for Computational Sciences (NICS) will again be competing for a next generation leadership class computing facility to be funded by the NSF. As with previous such competitions in which UTK and ORNL have participated over the past fifteen years, I expect ICL to be an important contributor to the UTK team once again. Second, to grow and improve the local and regional workforce, we must compete in this area; we must also become educational leaders in the rapidly evolving field of computational science. Our work with CITR to develop the Interdisciplinary Graduate Minor in Computational Science (IGMCS) initiated that effort at UTK; it is my hope that this year, we can make another substantial stride upward, taking computational science education at UTK to the next level.

So the stage is set for what promises to be a very rousing and rewarding year for ICL, both in its main areas of research and in its work to grow and strengthen UTK and its broader community. During these exciting times, I am grateful to our sponsors for their continued support of our efforts. My special thanks and congratulations go to the ICL staff and students for their skill, dedication, and tireless efforts to make ICL one of the best centers in the world for scientific computing technologies.

Jack Dongarra
Director, ICL

Claxton Building Home of ICL

INTRODUCTION

The Innovative Computing Laboratory is a large Computer Science research and development group situated in the heart of the University of Tennessee's Knoxville campus. ICL's mission is to establish and maintain the University of Tennessee as a world leader in advanced scientific and high performance computing through research, education, and collaboration.

ICL's founder, Dr. Jack Dongarra, established the lab in 1989 when he received a dual appointment as a Distinguished Professor at UTK and as a Distinguished Scientist at Oak Ridge National Laboratory. Since then, ICL has grown into an internationally recognized research laboratory, specializing in Numerical Linear Algebra, Performance Evaluation and Benchmarking, and Distributed Computing. The lab now employs nearly fifty researchers, students, and staff, and has earned many accolades, including four R&D100 awards.

23 YEARS OF **KEY ICL RESEARCH**

1989 PVM, Level 3 BLAS

1992 LAPACK, BLACS

1993 TOP500

1994 MPI

1995 ScaLAPACK

1997 NetSolve, ATLAS, RIB

1999 PAPI, HARNESS

2000 HPL

2002 FT-MPI, GCO

2003 LAPACK for Clusters, HPC Challenge

2006 Open MPI, FT-LA

2008 PLASMA, MAGMA

2009 Blackjack, MuMMI

2010 DPLASMA

2011 PULSAR

2012 PaRSEC

SELECTED **ICL ALUMNI**

1991-1992
Adam Beguelin
SENSR.NET

1995-1996
Sven Hammarling
NUMERICAL ALGORITHMS GROUP

2004-2007
Alfredo Buttari
CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE

1991-2001
Clint Whaley
UNIVERSITY OF TEXAS, SAN ANTONIO

1995-1998
Henri Casanova
UNIVERSITY OF HAWAII, MANOA

2003-2004
Edgar Gabriel
UNIVERSITY OF HOUSTON

2006-2008
Karl Fuerlinger
LUDWIG-MAXIMILIANS-UNIVERSITY MUNICH

1989-2001
Susan Blackford
MYRICOM

1994-2009
Keith Seymour
NATIONAL INSTITUTE FOR COMPUTATIONAL SCIENCES

1992-2005
Victor Eijkhout
UNIVERSITY OF TEXAS, AUSTIN

1992-1995
Jaeyoung Choi
SOONGSIL UNIVERSITY, KOREA

1996-2001
Dorian Arnold
UNIVERSITY OF NEW MEXICO

2004-2009
Haihang You
NATIONAL INSTITUTE FOR COMPUTATIONAL SCIENCES

1993-2001
Antoine Petitet
ESI GROUP, FRANCE

2003-2006
Julien Langou
UNIVERSITY OF COLORADO, DENVER

1989-1991
Ed Anderson
EPA

1996-2006
Graham Fagg
MICROSOFT

1991-1994
Roldan Pozo
NIST

1995-2001
Erich Strohmaier
LAWRENCE BERKELEY NATIONAL LABORATORY

2003-2006
Felix Wolf
JÜLICH SUPERCOMPUTING CENTRE

1992-1994
Richard Barrett
SANDIA NATIONAL LABORATORIES

1996-1997
Yves Robert
ENS-LYON, FRANCE

2009
Emmanuel Agullo
INRIA, FRANCE

1989-1990
Zhaojun Bai
UNIVERSITY OF CALIFORNIA, DAVIS

1995-1996
Greg Henry
INTEL

1999-2003
Sathish Vadhiyar
INDIAN INSTITUTE OF SCIENCE, INDIA

1990-1996
Bob Manchek
STRATUS TECHNOLOGIES

1999-2010
David Cronk
LOCKHEED-MARTIN

1990-1991
Robert van de Geijn
UNIVERSITY OF TEXAS, AUSTIN

1995-1996
Andy Cleary
AMAZON

2001-2006
Zizhong (Jeffrey) Chen
COLORADO SCHOOL OF MINES

2008-2011
Hatem Ltaief
KAUST, SAUDI ARABIA

1993-1994
Bernard Tourancheau
UNIVERSITY OF LYON, FRANCE

1997
Françoise Tisseur
MANCHESTER UNIVERSITY, ENGLAND

1994
Frederic Desprez
ENS-LYON, FRANCE

1999-2004
Ken Roche
UNIVERSITY OF WASHINGTON

1993-2012
Shirley Moore
UNIVERSITY OF TEXAS AT EL-PASO

1994-2004
Phil Mucci

1990-1996
Reed Wade
WETA DIGITAL

RESEARCH

Increased efforts to keep pace with the evolution in HPC hardware and software represent unique challenges that only a handful of enabling technology researchers are capable of addressing successfully. Our cutting-edge research efforts of the past have provided the foundation for addressing these challenges and serve as catalysts for success in our ever growing research portfolio. Our vision, our expertise, our determination, and our track record continue to position ICL as a leader in academic research.

What originally began more than 20 years ago as in-depth investigations of the numerical libraries that encode the use of linear algebra in software, has grown into an extensive research portfolio. We have evolved and expanded our research agenda to accommodate the aforementioned evolution of the HPC community, which includes a focus on algorithms and libraries for multicore and hybrid computing. As we have gained a solid understanding of the challenges presented in these domains, we have further expanded our research to include work in performance evaluation and benchmarking for high-end computers, as well as work in high performance parallel and distributed computing, with efforts focused on message passing and fault tolerance.

Demonstrating the range and diversity of our research, we will be engaged in more than 20 significant research projects during 2012-2013 across our main areas of focus.

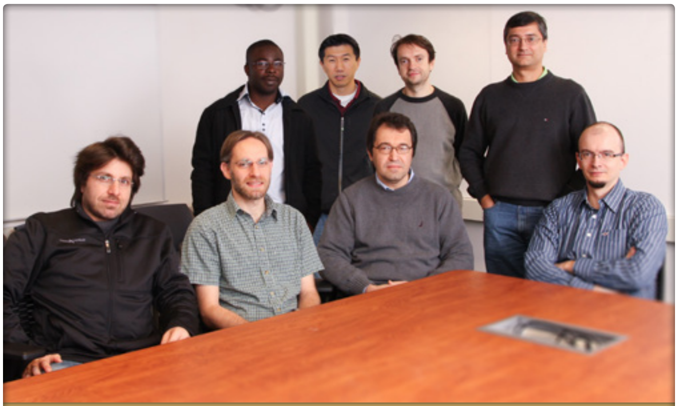

NUMERICAL LINEAR ALGEBRA

Numerical Linear Algebra algorithms and software form the backbone of many scientific applications in use today. With the ever-changing landscape of computer architectures, such as the massive increase in parallelism and the introduction of hybrid platforms utilizing both traditional CPUs and GPUs, these libraries must be revolutionized in order to achieve high performance and efficiency on these new hardware platforms. ICL has a long history of developing and standardizing these libraries in order to meet this demand, and we have multiple projects actively underway in this arena.

Numerical Linear Algebra Group
(Back Row) Simplice Donfack, Ichitaro Yamazaki, Piotr Luszczek, Asim YarKhan; (Front Row) Azzam Haidar, Mark Gates, Stanimire Tomov, Jakub Kurzak

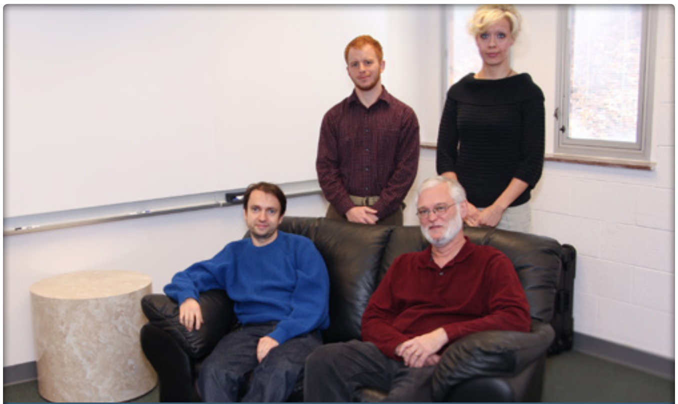

PERFORMANCE EVALUATION AND BENCHMARKING

Performance Evaluation and Benchmarking are vital to developing science and engineering applications that run efficiently in an HPC environment. ICL's Performance Evaluation tools allow programmers to see the correlation between the structure of source/object code and the efficiency of the mapping of that code to the underlying architecture. These correlations are important for performance tuning, compiler optimization, debugging, and finding and correcting performance bottlenecks. ICL's benchmark software is widely used to determine the performance profile of modern HPC machines, and has come to play an essential role in the purchasing and management of major computing infrastructure by government and industry around the world.

Performance Evaluation and Benchmarking Group
(Back Row) James Ralph, Heike McCraw;
(Front Row) Piotr Luszczek, Dan Terpstra

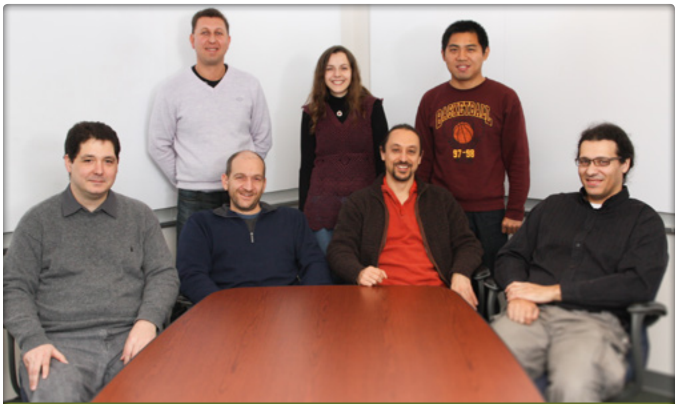

DISTRIBUTED COMPUTING

Distributed Computing is an integral part of the high performance computing landscape. As the number of cores, nodes, and other components in an HPC system continue to grow explosively, applications need runtime systems that can exploit all this parallelism. Moreover, the drastically lower meantime to failure of these components must be addressed with fault tolerant software and hardware, and the escalating communication traffic that they generate must be addressed with smarter and more efficient message passing standards and practices. Distributed Computing research at ICL has been a priority for nearly two decades, and the lab has several projects in that arena under active development.

Distributed Computing Group
(Back Row) Volodymyr Turchenko, Stephanie Moreaud, Erlin Yao;
(Front Row) Thomas Herault, George Bosilca, Anthony Danalis, Aurelien Bouteiller

DOD PETTT CREATE

High-Performance Numerical Libraries with Support for Low-Rank Matrix Computations

The main objective of the High-Performance Numerical Libraries with Support for Low-Rank Matrix Computations project is to reinforce a number of hybrid finite element/boundary integral analysis codes with high-performance numerical libraries that will enable them to efficiently use current and upcoming heterogeneous multi/many-core CPU and GPU-based architectures. The applications of interest use matrix compression techniques to reduce the complexity in both computer memory and computations. The compression techniques use adaptive cross approximation and SVD approaches that lead to low-rank matrix computations. Matrix factorizations and system solvers are similar to non-compressed methods except that the underlying matrix computations are performed in low-rank form.

ICL's work on this project focuses on the development of GPU algorithms, and their CUDA/OpenCL implementations for the required low-rank matrix computations, as well as the high level factorizations and solvers that will support matrix compression through the newly developed low-rank matrix computations. Of interest are factorizations and solvers for dense matrices, and multi-frontal methods for sparse matrices.

EASI

Extreme-scale Algorithms and Software Institute

The mission of the Extreme-scale Algorithms and Software Institute (EASI) is to close the performance gap between the peak capabilities of HPC hardware and the performance realized by high performance computing applications. To carry out this mission, the EASI project team develops architecture-aware algorithms and libraries, and the supporting runtime capabilities, to achieve scalable performance and resilience on heterogeneous architectures.

The project team includes personnel from ORNL, Sandia National Laboratories, University of Illinois, University of California Berkeley, and the University of Tennessee (ICL). ICL's efforts focus on providing components and services in a vertically integrated software stack, from low-level runtime process and thread scheduling to multicore aware library interfaces, multicore dense linear algebra, scalable iterative methods, and advanced parallel algorithms that break traditional parallelism bottlenecks.

FT-LA

Fault Tolerant Linear Algebra

The Fault Tolerant Linear Algebra (FT-LA) research effort is aimed at understanding and developing Algorithm Based Fault Tolerance (ABFT) into major dense linear algebra kernels. With distributed machines currently reaching up to 300,000 cores, fault-tolerance has never been so paramount. The scientific community has to tackle process failures from two directions: first, efficient middleware needs to be designed to detect failures, and second, the numerical applications have to be flexible enough to permit the recovery of the lost data structures.

At ICL, we have successfully developed Fault Tolerant MPI (FT-MPI) middleware and, more recently, an FT-LA library that will efficiently handle several process failures. The project team has also integrated FT-LA in the CIFTS (Coordinated Infrastructure for Fault Tolerant Systems) environment to provide better communication and fault management between the system's software components and scientific applications. Future work in this area involves the development of scalable fault-tolerant, one-sided (Cholesky, LU, and QR) and two-sided (Hessenberg, tri-diagonalization, and bi-diagonalization) factorizations, following the ABFT principles.

FIND OUT MORE AT http://icl.utk.edu/ft-la/

KEENELAND

Keeneland is a five-year, $12 million cyberinfrastructure project, awarded in 2009 under the NSF's Track 2D program, designed to bring emerging hardware architectures to the open science community. ICL is partnering with project leader Georgia Tech, as well as Oak Ridge National Laboratory, the National Institute for Computational Sciences, HP, and NVIDIA, to develop and deploy Keeneland's innovative and experimental system.

As part of our contribution, ICL performed education and outreach activities, developed numerical libraries to leverage the power of NVIDIA's CUDA-based GPUs used in the Keeneland machine, and teamed up with early adopters to map their applications to the Keeneland architecture. In 2012, the Keeneland Full Scale (KFS) system was accepted by the NSF and went into production. KFS is a ~615 teraflop/s, 264-node HP SL250 system with 528 Intel Sandy Bridge CPUs and 792 NVIDIA M2090 graphics processors, connected by an InfiniBand FDR network.

FIND OUT MORE AT http://keeneland.gatech.edu/

LAPACK

Linear Algebra PACKage

SCALAPACK

Scalable LAPACK

The Linear Algebra PACKage (LAPACK) and Scalable LAPACK (ScaLAPACK) are widely used libraries for efficiently solving dense linear algebra problems. ICL has been a major contributor to the development and maintenance of these two packages since their inception. LAPACK is sequential, relies on the BLAS library, and benefits from the multicore BLAS library. ScaLAPACK is parallel distributed and relies on BLAS, LAPACK, MPI, and BLACS libraries.

LAPACK 3.4.0 was released in November 2011 and includes LAPACKE, a native C interface for LAPACK, developed in collaboration with INTEL, which provides NAN check and automatic workspace allocation. LAPACK 3.4.0 also includes improved interface routines for QR factorization, new recursive QR factorization routines, and Communication-Avoiding QR sequential kernels. There were two new LAPACK releases in 2012 (3.4.1 and 3.4.2) for minor bug fixes, and the next major release is expected in January 2013. ScaLAPACK 2.0.0, which includes the MRRR algorithm and new Nonsymmetric Eigenvalue Problem routines, was released in November 2011. Two additional ScaLAPACK versions (2.0.1 and 2.0.2) were released in 2012 for minor bug fixes.

FIND OUT MORE AT http://www.netlib.org/lapack/
http://www.netlib.org/scalapack/

MAGMA

Matrix Algebra on GPU and Multicore Architectures

Matrix Algebra on GPU and Multicore Architectures (MAGMA) is a collection of next generation linear algebra (LA) libraries for heterogeneous architectures. The MAGMA package supports interfaces for current LA packages and standards, e.g., LAPACK and BLAS, to allow computational scientists to easily port any LA-reliant software components to heterogeneous architectures. MAGMA allows applications to fully exploit the power of current heterogeneous systems of multi/many-core CPUs and multi-GPUs/coprocessors to deliver the fastest possible time to accurate solution within given energy constraints.

MAGMA 1.3 features top performance and high accuracy LAPACK compliant routines for multicore CPUs enhanced with NVIDIA GPUs. The MAGMA 1.3 release includes more than 320 routines, covering one-sided dense matrix factorizations and solvers, two-sided factorizations and eigen/singular-value problem solvers, as well as a subset of highly optimized BLAS for GPUs. MAGMA provides multiple precision arithmetic support (S/D/C/Z, including mixed-precision). All algorithms are hybrid, using both multicore CPUs and GPUs. Support for AMD GPUs (OpenCL) and Intel Xeon Phi Coprocessors is provided correspondingly in the clMAGMA 1.0 and MAGMA MIC 0.3 releases.

FIND OUT MORE AT http://icl.utk.edu/magma/

PLASMA

Parallel Linear Algebra Software for Multicore Architectures

The Parallel Linear Algebra Software for Multicore Architectures (PLASMA) package, ICL's flagship project for multicore and many-core computing, is designed to deliver high performance from homogeneous multi-socket multicore systems by combining state-of-the-art solutions in algorithms, scheduling, and software engineering.

PLASMA provides routines for solving linear systems of equations and least square problems, including mixed-precision routines based on iterative refinement. PLASMA also provides routines for solving symmetric eigenvalue problems, generalized symmetric eigenvalue problems, and singular value problems. In addition, PLASMA contains routines for explicitly computing an inverse of a matrix, estimating matrix condition number, and rank-revealing QR factorization. PLASMA also includes a full set of Level 3 BLAS routines on matrices stored by tiles, and a set of routines to perform a wide range of matrix layout translations.

FIND OUT MORE AT http://icl.utk.edu/plasma/

PULSAR

Parallel Unified Linear Algebra with Systolic Arrays

The objective of the Parallel Unified Linear Algebra with Systolic Arrays (PULSAR) project is to address the challenges of extreme scale computing by applying the dataflow principles of Systolic Array architectures. PULSAR's solution is a Virtual Systolic Array (VSA) architecture, where multidimensional virtual systolic arrays are designed for various scientific workloads and successively mapped to the hardware architecture by a virtualization layer with a substantial runtime component.

Currently, PULSAR is focused on dense linear algebra codes, produced by mapping tile algorithms for matrix factorizations to VSAs, and then mapping those VSAs to the cores of a large distributed memory system. Initial results with a VSA implementation of the tile QR factorization show the huge potential of the approach. PULSAR's QR factorization of a 40Kx40K matrix scaled to 20K cores on the Kraken supercomputer at the National Institute for Computational Sciences (NICS), outperforming the vendor library by more than 3 times, for this particular problem.

FIND OUT MORE AT http://icl.utk.edu/pulsar/

QUARK

QUeuing And Runtime for Kernels

QUeuing And Runtime for Kernels (QUARK) provides a library that enables the dynamic execution of tasks with data dependencies in a multi-core, multi-socket, shared-memory environment. QUARK infers data dependencies and precedence constraints between tasks based on the way the data is used, and then executes the tasks in an asynchronous, dynamic fashion in order to achieve a high utilization of the available resources.

QUARK is designed to be easy to use, is intended to scale to large numbers of cores, and should enable the efficient expression and implementation of complex algorithms. The driving application behind the development of QUARK is the PLASMA linear algebra library, and the QUARK runtime contains several optimizations inspired by the algorithms in PLASMA.

FIND OUT MORE AT http://icl.utk.edu/quark/

FutureGrid

FutureGrid is a distributed computing infrastructure that uses the nation's high performance research networks to create a wide area test-bed for developing and testing novel approaches to parallel, grid, and cloud computing. In order to present a homogeneous resource to users, FutureGrid uses virtualization technology to create a cloud computing environment that applications can access and utilize in a uniform way. FutureGrid partners are deploying different kinds of high performance computing clusters at their sites and, at the same time, connecting them to XSEDE — the Extreme Science and Engineering Discovery Environment, funded by the NSF and based at NCSA. XSEDE, which is a follow-on to the TeraGrid project, is the NSF's newest national cyberinfrastructure for scientific research.

ICL is contributing to the project in the areas of performance measurement and application benchmarking, which pose very difficult problems for virtualized environments like FutureGrid. The recent release of PAPI 5.0, also known as PAPI-V, provides timing measurement for virtual and cloud environments, and has been deployed both on physical hardware for FutureGrid and in virtual space. For application benchmarking, ICL's Grid Benchmark Challenge effort applies the principles behind the HPC Challenge to assess virtualization overheads in controlled benchmarking scenarios.

FIND OUT MORE AT https://portal.futuregrid.org/

HPCC

HPC Challenge

The HPC Challenge (HPCC) benchmark suite is designed to assess the bounds of performance on many real-world applications for computational science at extreme scale. Included in the benchmark suite are tests for sustained floating point operations, memory bandwidth, rate of random memory updates, interconnect latency, and interconnect bandwidth. The main factors that differentiate the various components of the suite are the memory access patterns that, in a meaningful way, span the memory utilization space of temporal and spatial locality.

Each year, the HPCC Awards competition features contestants who submit performance numbers from the world's largest supercomputer installations, as well as alternative implementations that use a vast array of parallel programming environments. Results from the competition are announced at the annual Supercomputing Conference, and are available to the public to help track the progress of both the high-end computing arena and the commodity hardware segment.

FIND OUT MORE AT http://icl.utk.edu/hpcc/

HPL

High Performance LINPACK

The High Performance LINPACK (HPL) benchmark is a software package that solves a (random) dense linear system in double precision (64-bit) arithmetic on distributed-memory computers. Written in a portable ANSI C and requiring an MPI implementation as well as either the BLAS or VSIPL library, HPL is often one of the first programs to run on large computer installations, producing a result that can be submitted to the biannual TOP500 list of the world's fastest supercomputers.

HPL 2.1, released in 2012, includes several major bug fixes and accuracy enhancements based on user feedback. The major focus of HPL 2.1 is to improve the accuracy of reported benchmark results, and ensure scalability of the code on large supercomputer installations with hundreds of thousands of computational cores. The new version also features detailed time-of-run accounting to help with assessing power requirements at the time of execution, a metric which has been reported with TOP500 results since 2007 and is also highlighted on the Green500 list. In 2011, the LINPACK benchmark app for iOS achieved performance of over 1 gigaflop/s on an Apple iPad 2, with per-Watt performance easily beating supercomputing solutions, including the most power-efficient systems based on hardware accelerators.

FIND OUT MORE AT http://icl.utk.edu/hpl/

PERFORMANCE EVALUATION AND BENCHMARKING

MuMMI

Multiple Metrics Modeling Infrastructure

The Multiple Metrics Modeling Infrastructure (MuMMI) project is developing a framework to facilitate systematic measurement, modeling, and prediction of performance, power consumption, and performance-power tradeoffs for applications running on multicore systems. MuMMI combines UTK's PAPI hardware performance monitoring capabilities with Texas A&M's Prophesy performance modeling interface and Virginia Tech's Power-Pack power-performance measurement and analysis system.

PAPI has been integrated with Power-Pack and Prophesy, and performance and power consumption data have been collected for a range of benchmarks and applications running on multicore systems. Recent research activities at UTK have focused on developing PAPI hardware performance monitoring capabilities for GPU architectures, and on extending PAPI to measure power and energy consumption on multicore and GPU architectures.

FIND OUT MORE AT http://icl.utk.edu/mummi/

NSF PetaApps

Multiscale Software for Quantum Simulations in Nano Science and Technology

The Multiscale Software for Quantum Simulations in Nano Science and Technology project, led by North Carolina State University, is concerned with ab initio methods for computing the properties of materials and molecular structures. Petascale hardware enables ab initio calculations at unprecedented scale. The real-space multigrid (RMG) method being developed by this project uses a real-space mesh to represent the wavefunctions, the charge density, and the ionic pseudopotentials. The real-space formulation is advantageous for parallelization, since each processor can be assigned a region of space; and for convergence acceleration, since multiple length scales can be dealt with separately. Recent work focuses on implementing key algorithms on GPUs, and on overcoming memory and communication bottlenecks to allow RMG to scale to a hundred thousand cores, enabling simulations for thousands of atoms.

PAPI

Performance API

The Performance API (PAPI) has become the de facto standard within the HPC community for providing access to the hardware performance counters found on modern high performance computing systems. Provided as a linkable library or shared object, PAPI can be called directly in a user program, or used transparently through a variety of third party tools. Architecturally, PAPI provides simultaneous access to both on-processor and off-processor counters and sensors. PAPI continues to be ported to the architectures of greatest interest to the high performance computing community, including heterogeneous computing systems and virtual computing environments. Industry liaisons with Cray, IBM, NVIDIA, and others ensure seamless integration of PAPI with new architectures at or near their release. For example, close collaboration with IBM engineers allowed full PAPI support of the IBM Blue Gene/Q architecture when it was released.

Recent versions of PAPI support component technology, allowing PAPI to measure performance beyond just the CPU. Components are available for network counters, GPGPUs such as NVIDIA, system health monitoring, and most recently, power and energy measurement. Virtual PAPI, or PAPI-V—supported by NSF and VMware—was released in an initial implementation with PAPI 5.0. This version of PAPI will ultimately provide performance measurement standards in virtual environments, which are common in cloud computing. As power and energy consumption become an increasingly important aspect of high performance computing, PAPI is developing components to aid in their measurement. The current version of PAPI supports measuring energy consumption on Intel Sandy Bridge and Ivy Bridge processors through the RAPL (Running Average Power Limit) interface, and work is underway to support a variety of other methods for power and energy measurement. As the PAPI component ecosystem becomes more populated, performance tools that interface with PAPI automatically inherit the ability to measure these new data sources, usually with no changes at the user level. This provides a richer environment in which performance analysts can work.

FIND OUT MORE AT http://icl.utk.edu/papi/

PERFORMANCE EVALUATION AND BENCHMARKING

SUPER

Institute for Sustained Performance, Energy, and Resilience

The Institute for Sustained Performance, Energy, and Resilience (SUPER), led by the University of Southern California, has organized a broad-based project involving several universities and DOE laboratories with expertise in compilers, system tools, performance engineering, energy management, and resilience to ensure that DOE's computational scientists can successfully exploit the emerging generation of high performance computing (HPC) systems.

SUPER is extending performance modeling and autotuning technology on heterogeneous and Petascale computing systems, investigating application-level energy efficiency techniques, exploring resilience strategies for Petascale applications, and developing strategies that collectively optimize performance, energy efficiency, and resilience. UTK work focuses on performance measurement, power and energy measurements, and resilience techniques for hard and soft errors.

FIND OUT MORE AT http://super-scidac.org/

TOP500

Since 1993, a ranking of the top 500 fastest computers in the world has been compiled biannually with published results released in June and November. Each machine on the TOP500 is ranked based on performance results from running the computationally intensive High Performance LINPACK (HPL) benchmark developed by ICL.

While other benchmarks, including HPCC, have been developed to measure performance of HPC systems, the TOP500 still relies on the HPL benchmark and remains the de-facto ranking relied upon by commercial, industrial, government, and academic institutions. ICL continues to partner with NERSC/Lawrence Berkeley National Laboratory and the University of Mannheim, Germany to produce the rankings.

In November 2012 the Department of Energy's Titan supercomputer attained the #1 spot on the TOP500 list to become the fastest supercomputer on the planet. Titan, a Cray XK7 behemoth, is housed at the Oak Ridge National Laboratory, and achieved 17.59 petaflop/s on the HPL benchmark using 299,008 CPU cores and 18,688 GPU accelerators. Titan's raw performance is staggering, but it is also one of the most efficient machines on the TOP500 at 2,142.77 megaflop/s per watt.

FIND OUT MORE AT http://top500.org/

DISTRIBUTED COMPUTING

DPLASMA

Distributed Parallel Linear Algebra Software for Multicore Architectures

The Distributed Parallel Linear Algebra Software for Multicore Architectures (DPLASMA) package is the leading implementation of a dense linear algebra package for distributed heterogeneous systems. It is designed to deliver sustained performance for distributed systems where each node features multiple sockets of multicore processors, and if available, accelerators like GPUs or Intel Xeon Phi. DPLASMA achieves this objective through the state of the art PaRSEC runtime, by leveraging the Parallel Linear Algebra Software for Multicore Architectures (PLASMA) algorithms on the distributed memory realm.

In addition to traditional ScaLAPACK data distribution, DPLASMA provides interfaces for users to expose arbitrary tile distributions. The algorithms transparently operate on local data, or introduce implicit communications to resolve dependencies, removing the burden of initial data re-shuffle, and providing to the user a novel approach to address load balance.

FIND OUT MORE AT http://icl.utk.edu/parsec/

G8 ECS

Enabling Climate Simulation at Extreme Scale

The objective of the G8 Enabling Climate Simulation at Extreme Scale (ECS) project is to investigate how climate scientists can efficiently run climate simulations on future Exascale systems. Exascale supercomputers will appear in 2018–2020 featuring a hierarchical design, and will utilize hundreds of millions of computing cores. The numerical models of the physics, chemistry, and biology affecting the climate system need to be improved to run efficiently on these massive systems.

This project gathers top minds in climate research and computer science to focus on the challenges in resilience, performance, and scalability when running these simulations at extreme scale. The project team includes personnel from the University of Illinois at Urbana-Champaign, University of Tennessee (ICL), University of Victoria, German Research School for Simulation Sciences, INRIA, Barcelona Supercomputing Center, Tokyo Institute of Technology, and the University of Tsukuba. ICL's main role in the ECS project is moving widely used climate model code toward Exascale, starting with node level performance and scalability, and then application resilience.

GRIDPAC

Grid with Power-Aware Computing

Grid with Power-Aware Computing (GridPAC) is a middleware environment that will schedule multiple workflows across a distributed grid for system-wide optimization. GridPAC research focuses on supporting workflow execution using novel scheduling techniques on dynamic and heterogeneous resources. This project is a collaborative effort between the University of Texas Arlington, University of Florida, Virginia Tech, and the University of Tennessee, Knoxville.

At UTK, research has centered on the dynamic scheduling of DAG-structured workflows on a wide variety of computing resources, including multicore systems, systems utilizing multiple GPU resources, and grid-like distributed-memory, multicore systems using GPUs. The workflow applications are drawn from linear algebra, and thus they are likely to have a significant impact on multiple areas of scientific computing. GridPAC research is reflected in the PLASMA, QUARK, MAGMA, and PaRSEC projects.

FIND OUT MORE AT http://icl.utk.edu/gridpac/

NSF PetaApps

A Petaflop Cyberinfrastructure for Computing Free Energy Landscapes of Macro-and Biomolecular Systems

The objective of the Petaflop Cyberinfrastructure for Computing Free Energy Landscapes of Macro-and Biomolecular Systems project is to develop an Exascale cyberinfrastructure for the efficient calculation of free energy landscapes for complex macro- and bio-molecular systems. We then plan to demonstrate its effectiveness by applying it to two outstanding science problems: conformations of a linker protein in solution and self-assembly of lipids. In the context of this project, we have begun investigating, implementing, and evaluating different approaches to improve the scalability and the resilience of the cyberinfrastructure software computing for free energy landscapes. Our efforts are aimed at two axes of research. On one side, we are analyzing the current scalability of the existing framework and potential ways to drastically improve it. On the other side, we are investigating the software needs in order to achieve a reasonable level of resilience in this cyberinfrastructure.

Open MPI

The Open MPI Project is an open source Message Passing Interface (MPI) implementation that is developed and maintained by a consortium of academic, research, and industry partners. MPI primarily addresses the message-passing parallel programming model, in which data is moved from the address space of one process to that of another process through cooperative operations on each process. Open MPI integrates technologies and resources from several other projects (HARNESS/FT-MPI, LA-MPI, LAM/MPI, and PACX-MPI) in order to build the best MPI library available. A completely new MPI-2.2 compliant implementation, Open MPI offers advantages for system and software vendors, application developers, and computer science researchers.

ICL's efforts in the context of Open MPI have significantly improved its scalability, performance on many-core environments, and architecture aware capabilities, such as adaptive shared memory behaviors and dynamic collective selection, making it ready for the next generation Exascale challenges.

FIND OUT MORE AT http://icl.cs.utk.edu/open-mpi/

PaRSEC

Parallel Runtime Scheduling and Execution Controller

The Parallel Runtime Scheduling and Execution Controller (PaRSEC) is a generic framework for architecture aware scheduling and management of micro-tasks on distributed many-core heterogeneous architectures. Applications we consider are expressed as a Direct Acyclic Graph (DAG) of tasks, with edges designating data dependencies. DAGs are represented in a compact problem-size independent format that can be queried to discover data dependencies in a totally distributed fashion, a drastic shift from today's programming models, which are based on sequential flow of execution.

PaRSEC assigns computation threads to the cores, overlaps communications and computations, and uses a dynamic, fully-distributed scheduler based on architectural features such as NUMA nodes and algorithmic features such as data reuse. PaRSEC includes a set of tools to generate the DAGs and integrate them in legacy codes, a runtime library to schedule the micro-tasks on heterogeneous resources, and tools to evaluate and visualize the efficiency of the scheduling. Many dense and sparse linear algebra extensions have been implemented, and computational kernels have been re-implemented using PaRSEC, enabling better performance on distributed many-core systems.

FIND OUT MORE AT http://icl.utk.edu/parsec/

ULFM

User Level Failure Mitigation

User Level Failure Mitigation (ULFM) is a set of new interfaces for MPI that enables Message Passing programs to restore MPI functionality affected by process failures. The MPI implementation is spared the expense of internally taking protective and corrective actions against failures. Instead, it reports operations whose completions were rendered impossible by failures.

Using the constructs defined by ULFM, applications and libraries drive the recovery of the MPI state. Consistency issues resulting from failures are addressed according to an application's needs and the recovery actions are limited to the necessary MPI communication objects. Therefore, the recovery scheme is more efficient than a generic, automatic recovery technique, and can achieve both goals of enabling applications to resume communication after failure and maintaining extreme communication performance outside of recovery periods.

FIND OUT MORE AT http://fault-tolerance.org/

PUBLICATIONS

Abdelfattah, A., Dongarra, J., Keyes, D., Ltaief, H. "**Optimizing Memory-Bound Numerical Kernels on GPU Hardware Accelerators**," *VECPAR 2012*, Kobe, Japan, July, 2012.

Agullo, E., Bosilca, G., Castagnède, C., Dongarra, J., Ltaief, H., Tomov, S. "**Matrices Over Runtime Systems at Exascale**," *Supercomputing '12 (poster)*, Salt Lake City, Utah, November, 2012.

Anzt, H., Dongarra, J., Heuveline, V. "**Weighted Block-Asynchronous Relaxation for GPU-Accelerated Systems**," *SIAM Journal on Computing (submitted)*, March, 2012.

Anzt, H., Luszczek, P., Dongarra, J., Heuveline, V. "**GPU-Accelerated Asynchronous Error Correction for Mixed Precision Iterative Refinement**," *EuroPar 2012 (also LAWN 260)*, Rhodes Island, Greece, August, 2012.

Anzt, H., Tomov, S., Dongarra, J., Heuveline, V. "**A Block-Asynchronous Relaxation Method for Graphics Processing Units**," *Journal of Parallel and Distributed Computing (submitted)*, October, 2012.

Anzt, H., Tomov, S., Dongarra, J., Heuveline, V. "**Weighted Block-Asynchronous Iteration on GPU-Accelerated Systems**," *Tenth International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms (Best Paper Award)*, Rhodes Island, Greece, August, 2012.

Anzt, H., Tomov, S., Gates, M., Dongarra, J., Heuveline, V. "**Block-asynchronous Multigrid Smoothers for GPU-accelerated Systems**," *ICCS 2012*, Omaha, NE, June, 2012.

Baboulin, M., Becker, D., Bosilca, G., Danalis, A., Dongarra, J. "**An efficient distributed randomized solver with application to large dense linear systems**," *ICL Technical Report*, ICL-UT-12-02, July 11, 2012.

Baboulin, M., Becker, D., Dongarra, J. "**A Parallel Tiled Solver for Symmetric Indefinite Systems On Multicore Architectures**," *IPDPS 2012*, Shanghai, China, May, 2012.

Baboulin, M., Donfack, S., Dongarra, J., Grigori, L., Remi, A., Tomov, S. "**A class of communication-avoiding algorithms for solving general dense linear systems on CPU/GPU parallel machines**," *Proc. of the International Conference on Computational Science (ICCS)*, 9, 17-26, June, 2012.

Baboulin, M., Dongarra, J., Herrmann, J., Tomov, S. "**Accelerating Linear System Solutions Using Randomization Techniques**," *ACM Transactions on Mathematical Software (accepted) (also LAWN 246)*, March, 2012.

Becker, D., Baboulin, M., Dongarra, J. "**Reducing the Amount of Pivoting in Symmetric Indefinite Systems**," *Parallel Processing and Applied Mathematics, Lecture Notes in Computer Science (PPAM 2011)*, Wyrzykowski, R., Dongarra, J., Karczewski, K., Wasniewski, J. eds. Springer-Verlag Berlin Heidelberg, Vol. 7203, pp. 133-142, 2012.

Bland, W. "**Enabling Application Resilience With and Without the MPI Standard**," *11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, Ottawa, Canada, May, 2012.

Bland, W., Bosilca, G., Bouteiller, A., Herault, T., Dongarra, J. "**A Proposal for User-Level Failure Mitigation in the MPI-3 Standard**," *University of Tennessee Electrical Engineering and Computer Science Technical Report*, UT-CS-12-693, February 24, 2012.

Bland, W., Bouteiller, A., Herault, T., Hursey, J., Bosilca, G., Dongarra, J. "**An Evaluation of User-Level Failure Mitigation Support in MPI**," *Proceedings of Recent Advances in Message Passing Interface - 19th European MPI Users' Group Meeting, EuroMPI 2012*, Springer, Vienna, Austria, September 23 - 26, 2012.

Bland, W., Du, P., Bouteiller, A., Herault, T., Bosilca, G., Dongarra, J. "**A Checkpoint-on-Failure Protocol for Algorithm-Based Recovery in Standard MPI**," *18th International European Conference on Parallel and Distributed Computing (Euro-Par 2012)*, Christos Kaklamanis, Theodore Papatheodorou and Paul Spirakis eds. Springer-Verlag, Rhodes, Greece, August 27-31, 2012.

Bland, W., Du, P., Bouteiller, A., Herault, T., Bosilca, G., Dongarra, J. "**Extending the Scope of the Checkpoint-on-Failure Protocol for Forward Recovery in Standard MPI**," *University of Tennessee Computer Science Technical Report*, UT-CS-12-702, 2012.

Bosilca, G., Bouteiller, A., Brunet, E., Cappello, F., Dongarra, J., Guermouche, A., Herault, T., Robert, Y., Vivien, F., Zaidouni, D. "**Unified Model for Assessing Checkpointing Protocols at Extreme-Scale**," *University of Tennessee Computer Science Technical Report (also LAWN 269)*, UT-CS-12-697, June, 2012.

Bosilca, G., Bouteiller, A., Danalis, A., Herault, T., Lemarinier, P., Dongarra, J. "**DAGuE: A generic distributed DAG Engine for High Performance Computing**," *Parallel Computing*, T. Hoefler eds. Elsevier, Vol. 38, No 1-2, 27-51, 2012.

Bosilca, G., Bouteiller, A., Danalis, A., Herault, T., Luszczek, P., Dongarra, J. "**Dense Linear Algebra on Distributed Heterogeneous Hardware with a Symbolic DAG Approach**," *Scalable Computing and Communications: Theory and Practice (to appear)*, Khan, S., Wang, L., Zomaya, A. eds. John Wiley & Sons, January, 2012.

Bosilca, G., Dongarra, J., Ltaief, H. "**Power Profiling of Cholesky and QR Factorizations on Distributed Memory Systems**," *Third International Conference on Energy-Aware High Performance Computing*, Hamburg, Germany, September, 2012.

Bouteiller, A., Herault, T., Bosilca, G., Dongarra, J. "**Correlated Set Coordination in Fault Tolerant Message Logging Protocols**," *Concurrency and Computation: Practice and Experience (accepted)*, March, 2012.

Danalis, A., Bouteiller, A., Bosilca, G., Dongarra, J., Herault, T. "**From Serial Loops to Parallel Execution on Distributed Systems**," *PPoPP 2012 (submitted)*, New Orleans, LA, February, 2012.

Danalis, A., Luszczek, P., Marin, G., Vetter, J., Dongarra, J. "**BlackjackBench: Portable Hardware Characterization with Automated Results Analysis**," *Computer Journal (submitted)*, April, 2012.

Donfack, S., Tomov, S., Dongarra, J. "**Performance evaluation of LU factorization through hardware counter measurements**," *University of Tennessee Computer Science Technical Report*, UT-CS-12-700, October, 2012.

Dong, T., Kolev, T., Rieben, R., Dobrev, V., Tomov, S., Dongarra, J. "**Acceleration of the BLAST Hydro Code on GPU**," *Supercomputing '12 (poster)*, Salt Lake City, Utah, November, 2012.

Dongarra, J., Faverge, M., Herault, T., Langou, J., Robert, Y. "**Hierarchical QR factorization algorithms for multi-core cluster systems**," *IPDPS 2012, the 26th IEEE International Parallel and Distributed Processing Symposium, IEEE Computer Society Press*, Shanghai, China, May 21-25, 2012.

Dongarra, J., Kurzak, J., Luszczek, P., Tomov, S. "**Dense Linear Algebra on Accelerated Multicore Hardware**," *High Performance Scientific Computing: Algorithms and Applications*, Berry, M., et al. eds. Springer-Verlag, London, UK, 2012.

Dongarra, J., Ltaief, H., Luszczek, P., Weaver, V. "**Energy Footprint of Advanced Dense Numerical Linear Algebra using Tile Algorithms on Multicore Architecture**," *The 2nd International Conference on Cloud and Green Computing (submitted)*, Xiangtan, Hunan, China, November, 2012.

Evidence of our research and our contributions to the HPC community might be best exemplified by the numerous publications we produce every year. Here is a listing of our most recent papers, including journal articles, book chapters, and conference proceedings. Many of these are available for download from our website.

Dongarra, J., Luszczek, P. "**HPC Challenge: Design, History, and Implementation Highlights**," *On the Road to Exascale Computing: Contemporary Architectures in High Performance Computing (to appear)*, Vetter, J. eds. Chapman & Hall/CRC Press, 2012.

Dongarra, J., van der Steen, A.J. "**High Performance Computing Systems: Status and Outlook**," *Acta Numerica*, Cambridge University Press, Cambridge, UK, Vol. 21, pp. 379-474, May, 2012.

Du, P., Bouteiller, A., Bosilca, G., Herault, T., Dongarra, J. "**Algorithm-Based Fault Tolerance for Dense Matrix Factorization**," *Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2012*, J. Ramanujam, P. Sadayappan eds. ACM, New Orleans, LA, USA, 225-234, February 25-29, 2012.

Du, P., Luszczek, P., Dongarra, J. "**High Performance Dense Linear System Solver with Resilience to Multiple Soft Errors**," *ICCS 2012*, Omaha, NE, June, 2012.

Du, P., Tomov, S., Dongarra, J. "**Providing GPU Capability to LU and QR within the ScaLAPACK Framework**," *University of Tennessee Computer Science Technical Report (also LAWN 272)*, UT-CS-12-699, September, 2012.

Du, P., Weber, R., Luszczek, P., Tomov, S., Peterson, G., Dongarra, J. "**From CUDA to OpenCL: Towards a Performance-portable Solution for Multi-platform GPU Programming**," *Parallel Computing*, Vol. 38, No. 8, pp. 391-407, August, 2012.

Haidar, A., Ltaief, H., Dongarra, J. "**Toward High Performance Divide and Conquer Eigensolver for Dense Symmetric Matrices**," *SIAM Journal on Scientific Computing (Accepted)*, July, 2012.

Haidar, A., Ltaief, H., Luszczek, P., Dongarra, J. "**A Comprehensive Study of Task Coalescing for Selecting Parallelism Granularity in a Two-Stage Bidiagonal Reduction**," *IPDPS 2012*, Shanghai, China, May, 2012.

Johnson, M., McCraw, H., Moore, S., Mucci, P., Nelson, J., Terpstra, D., Weaver, V., Mohan, T. "**PAPI-V: Performance Monitoring for Virtual Machines**," *CloudTech-HPC 2012*, Pittsburgh, PA, September 10-13, 2012.

Kasichayanula, K., Terpstra, D., Luszczek, P., Tomov, S., Moore, S., Peterson, G. "**Power Aware Computing on GPUs**," *SAAHPC '12 (Best Paper Award)*, Argonne, IL, July 10-11, 2012.

Kurzak, J., Luszczek, P., Dongarra, J. "**LU Factorization with Partial Pivoting for a Multicore System with Accelerators**," *IEEE Transactions on Parallel and Distributed Computing (accepted)*, August, 2012.

Kurzak, J., Luszczek, P., Faverge, M., Dongarra, J. "**Programming the LU Factorization for a Multicore System with Accelerators**," *Proceedings of VECPAR'12*, Kobe, Japan, April, 2012.

Kurzak, J., Luszczek, P., Gates, M., Yamazaki, I., Dongarra, J. "**Virtual Systolic Array for QR Decomposition**," *IPDPS 2013 (submitted)*, Boston, MA, October, 2012.

Kurzak, J., Luszczek, P., Tomov, S., Dongarra, J. "**Preliminary Results of Autotuning GEMM Kernels for the NVIDIA Kepler Architecture**," *LAWN 267*, 2012.

Kurzak, J., Nath, R., Du, P., Dongarra, J. "**An Implementation of the Tile QR Factorization for a GPU and Multiple CPUs**," *Applied Parallel and Scientific Computing*, Lristjan Jonasson eds. Vol. 7133, pp. 248-257, 2012.

Langou, J., Hoffman, B., King, B. "**How LAPACK library enables Microsoft Visual Studio support with CMake and LAPACKE**," *University of Tennessee Computer Science Technical Report (also LAWN 270)*, UT-CS-12-698, July, 2012.

Ltaief, H., Luszczek, P., Dongarra, J. "**Enhancing Parallelism of Tile Bidiagonal Transformation on Multicore Architectures using Tree Reduction**," *Lecture Notes in Computer Science*, Vol. 7203, pp. 661-670, September, 2012.

Luszczek, P., Dongarra, J. "**Anatomy of a Globally Recursive Embedded LINPACK Benchmark**," *2012 IEEE High Performance Extreme Computing Conference*, Waltham, MA, September, 2012.

Luszczek, P., Kurzak, J., Dongarra, J. "**Looking Back at Dense Linear Algebra Software**," *Perspectives on Parallel and Distributed Processing: Looking Back and What's Ahead (to appear)*, Prasanna, V., Robert, Y., Stenström, P. eds. 2012.

Ma, T., Bosilca, G., Bouteiller, A., Dongarra, J. "**HierKNEM: An Adaptive Framework for Kernel-Assisted and Topology-Aware Collective Communications on Many-core Clusters**," *IPDPS 2012 (Best Paper Award)*, Shanghai, China, May, 2012.

McCraw, H. "**Performance Counter Monitoring for the Blue Gene/Q Architecture**," *University of Tennessee Computer Science Technical Report*, ICL-UT-12-01, 2012.

"**Parallel Processing and Applied Mathematics, 9th International Conference, PPAM 2011**," *Lecture Notes in Computer Science*, Wyrzykowski, R., Dongarra, J., Karczewski, K., Wasniewski, J. eds. Torun, Poland, Vol. 7203, 2012.

"**Recent Advances in the Message Passing Interface: 19th European MPI Users' Group Meeting, EuroMPI 2012**," *Lecture Notes in Computer Science*, Larsson Träff, J., Benkner, S., Dongarra, J. eds. Vienna, Austria, Vol. 7490, 2012.

Solcà, R., Haidar, A., Tomov, S., Dongarra, J., Schulthess, T. "**A Novel Hybrid CPU-GPU Generalized Eigensolver for Electronic Structure Calculations Based on Fine Grained Memory Aware Tasks**," *Supercomputing '12 (poster)*, Salt Lake City, Utah, November, 2012.

Song, F., Dongarra, J. "**A Scalable Framework for Heterogeneous GPU-Based Clusters**," *The 24th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2012)*, ACM, Pittsburgh, PA, USA, June 25, 2012.

Song, F., Tomov, S., Dongarra, J. "**Enabling and Scaling Matrix Computations on Heterogeneous Multi-Core and Multi-GPU Systems**," *26th ACM International Conference on Supercomputing (ICS 2012)*, ACM, San Servolo Island, Venice, Italy, June, 2012.

Vomel, C., Tomov, S., Dongarra, J. "**Divide and Conquer on Hybrid GPU-Accelerated Multicore Systems**," *SIAM Journal on Scientific Computing, 34 (2)*, C70-C82, April 12, 2012.

Weaver, V., Johnson, M., Kasichayanula, K., Ralph, J., Luszczek, P., Terpstra, D., Moore, S. "**Measuring Energy and Power with PAPI**," *International Workshop on Power-Aware Systems and Architectures*, Pittsburgh, PA, September 10, 2012.

Yamazaki, I., Tomov, S., Dongarra, J. "**One-sided dense matrix factorizations on a multicore with multiple GPU accelerators**," *The International Conference on Computational Science (ICCS)*, June 4, 2012.

EVENTS

Every year, various members of our research staff regularly attend national and international conferences, workshops, and seminars. These meetings provide opportunities to present our research, share our knowledge, and exchange ideas with leading computational science researchers from around the world. Showing the work we have done and participating in the intellectual life of the scientific community in this way is an essential part of the research process.

JANUARY 9-11
MPI Forum
San Jose, CA

JANUARY 23-27
Third UiO-MSU-ORNL-UT School on Topics in Nuclear Physics
Oak Ridge, TN

FEBRUARY 5-8
Workshop on Electronic Structure Calculation Methods on Accelerators
Oak Ridge, TN

FEBRUARY 6-8
An Event Apart 2012
Atlanta, GA

FEBRUARY 15-17
SIAM Conference on Parallel Processing for Scientific Computing
Savannah, GA

FEBRUARY 20-21
Keeneland Workshop
Atlanta, GA

FEBRUARY 25-29
17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming
New Orleans, LA

FEBRUARY 28 - MARCH 1
DOE CREATE Developers Review Savannah, GA

MARCH 5-7
MPI Forum
Chicago, IL

MARCH 6-7
G8 Workshop
Aachen, Germany

MARCH 12
PetaApps Annual Meeting
Chattanooga, TN

MARCH 19-21
Early Science Program Code for Q Workshop
Argonne, IL

MARCH 28-30
Accelerating Computational Science Symposium
Washington, DC

MARCH 29-30
SUPER All-hands Meeting
Chapel Hill, NC

APRIL 2-6
International Young Scientists Conference 2012 Amsterdam, Netherlands

APRIL 4-5
ASEArch Workshop - Dynamic Execution and AutoTuning
Oxford, United Kingdom

APRIL 11-13
International Exascale Software Project
Kobe, Japan

APRIL 13
Stanford University ICME CUDA Day
Palo Alto, CA

APRIL 25-27
2012 Joint JICS/GRS Workshop on Large-Scale Computer Simulation
Oak Ridge, TN

APRIL 27-30
Scalable Hierarchical Algorithms for Extreme Computing
Jeddah, Saudi Arabia

APRIL 30 - MAY 2
Early Science Program Code for Q Workshop
Chicago, IL

MAY 7-10
ATIP/A*CRC Workshop on Accelerator Technologies for High Performance Computing
Biopolis, Singapore

MAY 13-16
CCGrid 2012
Ottawa, Ontario

MAY 14-17
GPU Technology Conference 2012
San Jose, CA

MAY 14-18
SciDOMP 2012
Toronto, Ontario

MAY 21-25
26th IEEE International Parallel & Distributed Processing Symposium
Shanghai, China

MAY 28-30
MPI Forum
Tsukuba, Japan

JUNE 4-6
International Conference on Computational Science (ICCS)
Omaha, NE

JUNE 6-8
OMPI Developers Meeting
San Jose, CA

JUNE 10-13
PARA 2012
Helsinki, Finland

JUNE 11-14
AMD Fusion Developer Summit
Bellevue, WA

JUNE 15-16
Codestock 2012
Knoxville, TN

JUNE 17-21
International Supercomputing Conference
Hamburg, Germany

JUNE 18-22
SIAM Conference on Applied Linear Algebra
Valencia, Spain

JUNE 22
Erlangen International High-End-Computing Symposium
Erlangen, Germany

JUNE 25-29
International Conference on Supercomputing (ICS 2012)
Venice, Italy

JUNE 25-27
ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2012) Pittsburgh, PA

JUNE 25-28
European-U.S. Summer School on HPC Challenges in Computational Sciences
Dublin, Ireland

JUNE 26-29
CScADS Performance Tools Workshop
Salt Lake City, UT

JUNE 28-30
7th Scheduling for Large Scale Systems Workshop
Pittsburgh, PA

JULY 16-20
XSEDE 2012
Chicago, IL

JULY 18-20
VECPAR 2012
Kobe, Japan

JULY 20
MuMMI Team Meeting
Chicago, IL

JULY 30 - AUGUST 3
MORSE Associate Team Visit
Bordeaux, France

AUGUST 13-15
CScADS Summer 2012 Workshop Snowbird, UT

AUGUST 14
PetaApps Meeting
Nashville, TN

AUGUST 27-31
International European Conference on Parallel and Distributed Computing (Euro-Par)
Rhodes Island, Greece

AUGUST 29
SciDB Workshop
Boston, MA

SEPTEMBER 4-7
2012 Smoky Mountains Computational Sciences and Engineering Conference
Gatlinburg, TN

SEPTEMBER 7
41st SPEEDUP Workshop on High-Performance Computing
Zurich, Switzerland

SEPTEMBER 10-13
International Conference on Parallel Processing
Pittsburgh, PA

SEPTEMBER 10-12
IEEE High Performance Embedded Computing Conference
Waltham, MA

SEPTEMBER 10-12
SciDAC 3 PI Meeting
Rockville / Washington DC, MD

SEPTEMBER 11-14
Clusters, Clouds, and Data for Scientific Computing (CCDSC)
Lyon, France

SEPTEMBER 18
SBIR/STTR Programs Overview Workshop
Oak Ridge, TN

SEPTEMBER 20-21
MPI Forum
Vienna, Austria

SEPTEMBER 24-26
EuroMPI
Vienna, Austria

SEPTEMBER 26-27
VMware Academic Research Symposium
Cambridge, MA

SEPTEMBER 26-27
SUPER All-hands Meeting
Argonne, IL

OCTOBER 3-5
CASC Fall Meeting
Alexandria, VA

OCTOBER 9-11
International Computational Mechanics Symposium 2012
Kobe, Japan

OCTOBER 28-29
IIT Computer Science Reunion 2012
Chicago, IL

NOVEMBER 1
Virginia Tech CS Distinguished Lecture
Blacksburg, VA

NOVEMBER 10-16
Supercomputing 2012
Salt Lake City, UT

NOVEMBER 28 - DECEMBER 2
The First International Mega-Conference on Green and Smart Technology
Jeju, Korea

DECEMBER 1-3
Saudi Arabian High Performance Computing Users' Group Conference
Thuwal, Saudi Arabia

DECEMBER 4
Seminar Université Paris-Sud
Orsay, France

DECEMBER 14
AFOSR Kickoff Meeting
Des Moines, IA

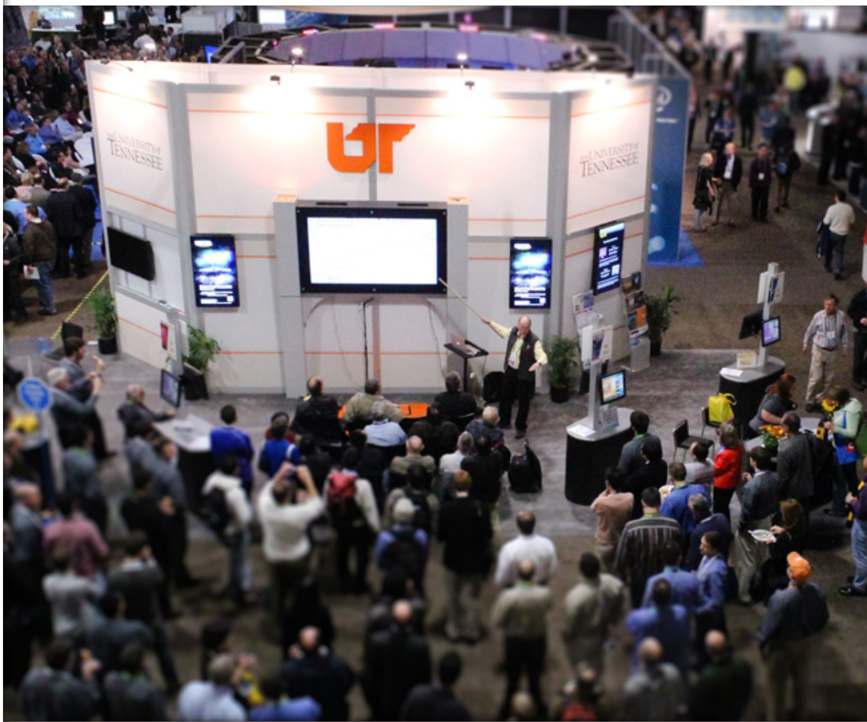

SC12

The annual ACM/IEEE Supercomputing Conference (SC), established in 1988, is a staple of ICL's November itinerary. SC is vital to the growth and evolution of high performance computing in the United States because it is the only US event that elicits substantial participation from all segments of the HPC community, including hundreds of users, developers, vendors, research institutions, and representatives of governmental funding agencies. Such a talent-rich gathering enables participants to discuss challenges, share innovations, and coordinate relationships and collaborations with some of the best minds in scientific and high performance computing.

ICL routinely has a significant presence at SC, with faculty, research staff, and students giving talks, presenting papers, and leading "Birds-of-a-Feather" sessions; SC12 was no exception. ICL also had a new venue this year in the form of the University of Tennessee's first ever SC booth. The booth, which was organized and led by the National Institute for Computational Sciences, was visually designed with the help of ICL/CITR staff, manned with support from ICL researchers attending SC, and featured the lab's research projects in the booth's kiosks. ICL's director and UT Distinguished Professor, Jack Dongarra, drew a large crowd to the booth during the exhibition's Gala Opening when he gave his annual talk presenting the November 2012 TOP500 list.

PARTNERSHIPS

Since 1989, ICL has fostered relationships with many other academic institutions and research centers. We have also aggressively sought to build lasting, collaborative partnerships with HPC vendors and industry research leaders, both here and abroad. Together with these partners, we have built a strong portfolio of shared resources, both material and intellectual. In this section, we recognize many of the partners and collaborators that we have worked with over the years, most of whom we are still actively involved with.

Vendors and industry research leaders are an integral part of our partnerships, contributing significantly to our efforts to be a world leader in computational science research. Many have utilized our work, including our linear algebra libraries and performance analysis tools. As a result of these exchanges, we maintain close working relationships with many industry leaders.

Our relationships with academic and government research institutions also play a pivotal role in our success. By exchanging ideas, expertise, and personnel, ICL becomes more dynamic with each new collaboration. Our lab routinely develops relationships with researchers whose primary focus is on other scientific disciplines, such as biology, chemistry, and physics, which makes many of our collaborations truly multidisciplinary.

COLLABORATING INTERNATIONAL INSTITUTIONS

Barcelona Supercomputing Center
Barcelona, Spain

Central Institute for Applied Mathematics
Jülich, Germany

Consiglio Nazionale delle Ricerche
Rome, Italy

Danish Computing Center for Research and Education
Lyngby, Denmark

Doshisha University
Kyoto, Japan

École Normale Supérieure de Lyon
Lyon, France

École Polytechnique Federale de Lausanne
Lausanne, Switzerland

European Centre for Research and Advanced Training in Scientific Computing
Toulouse, France

European Exascale Software Initiative
European Union

Forschungszentrum Jülich
Jülich, Germany

High Performance Computing Center Stuttgart
Stuttgart, Germany

INRIA
France

Institut ETH Zentrum
Zurich, Switzerland

Kasetsart University
Bangkok, Thailand

King Abdullah University of Science and Technology
Saudi Arabia

Parallel and HPC Application Software Exchange
Tsukuba, Japan

Rutherford Appleton Laboratory
Oxford, England

Soongsil University
Seoul, South Korea

Technische Universitaet Wien
Vienna, Austria

Tokyo Institute of Technology
Tokyo, Japan

Universität Mannheim
Mannheim, Germany

Université Claude Bernard de Lyon
Lyon, France

University of Manchester
Manchester, England

University of Umeå
Umeå, Sweden

VI-HPS

In 2007, ICL became part of a collaboration for HPC research called the Virtual Institute - High Productivity Supercomputing (VI-HPS), whose mission is "to improve the quality and accelerate the development process of complex simulation programs in science and engineering that are being designed for the most advanced parallel computer systems." The institute, comprised of institutions in Europe and the US, including ICL, unites some of the brightest minds in HPC research who are committed to helping engineers and domain scientists become more efficient and effective users of HPC applications.

Over the years, ICL's membership and contributions have proven valuable to the success of the institute. ICL's PAPI project provides access to hardware performance counters and other performance metrics that are leveraged by a range of other performance tools developed by VI-HPS partners, including TAU, Scalasca, Vampir, and others. ICL researchers have also participated in workshops and tutorials showcasing performance tools, including PAPI, in a variety of international venues. As one of the 11 VI-HPS partners, ICL anticipates a continuing fruitful collaboration with both industry and research institutions in furthering the optimization of performance in High Productivity Supercomputing.

VI-HPS PARTNERS

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Forschungszentrum Jülich
Jülich Supercomputing Centre

German Research School for Simulation Sciences
Laboratory of Parallel Programming

Lawrence Livermore National Laboratory
Computation Directorate

RWTH Aachen University
Center for Computing and Communication

Technische Universität Dresden
Center for Information Services and High Performance Computing

Technische Universität München
Chair for Computer Architecture

Université de Versailles St-Quentin-en-Yvelines
LRC ITACA

University of Oregon
Performance Research Laboratory

University of Stuttgart
High Performance Computing Centre

University of Tennessee
Innovative Computing Laboratory

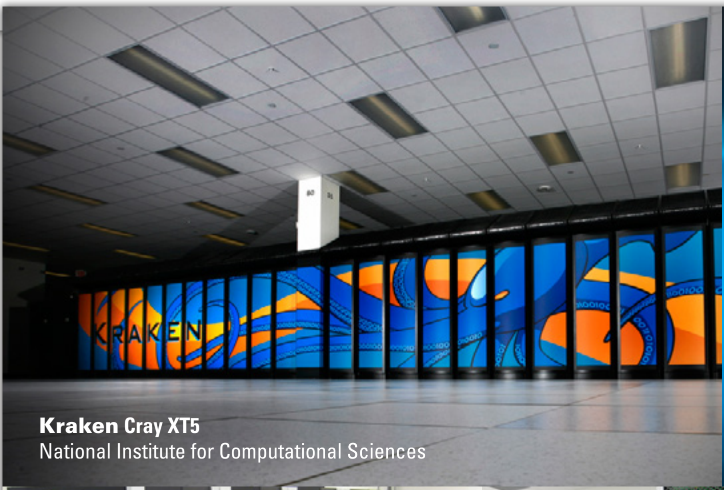

Kraken Cray XT5
National Institute for Computational Sciences

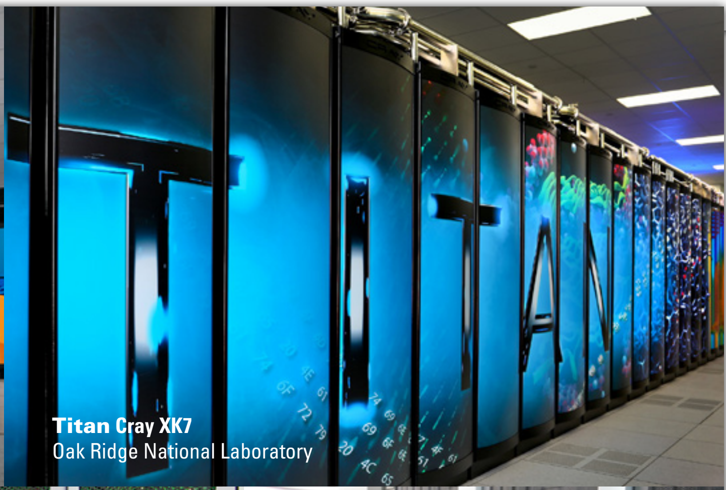
Titan Cray XK7
Oak Ridge National Laboratory


Newton Cluster
University of Tennessee

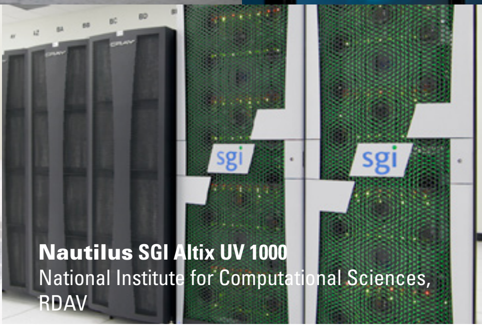

Nautilus SGI Altix UV 1000
National Institute for Computational Sciences, RDAV


Keeneland System
Georgia Tech

COMPUTATIONAL RESOURCES

As the new GPU hybrid computing paradigm leads the evolution of computational hardware into Petascale computing, computing architectures are increasingly changing. However, the programming tools, applications, and algorithms that form the backbone of the ever growing need for greater performance are equally as important. Such myriad hardware/software configurations present unique challenges that require testing and development of applications that are often unique to the platform on which they reside. For this reason, it is imperative that we have access to a wide range of computing resources in order to conduct our cutting-edge research.

CAMPUS RESOURCES

To meet these challenges, ICL has access to multiple state-of-the-art heterogeneous systems in house. In fact, our hardware industry partners, including AMD, ARM, Intel, and NVIDIA, provide us with bleeding edge hardware resources (often under NDA and prior to public release) which we use to upgrade and maintain the lab's infrastructure. Our research staff also has access to other campus resources, including UTK's Newton Cluster.

LOCAL RESOURCES

ICL has access to many local resources in East Tennessee to help keep us at the forefront of enabling technology research, including some machines that are regularly found on the TOP500 list of the world's fastest supercomputers. The recent modernization of the DOE's National Center for Computational Sciences (NCCS), just 30 minutes away at the Oak Ridge National Laboratory (ORNL), has enabled us to leverage our ORNL collaborations to take advantage of what has become one of the world's fastest scientific computing facilities.

ORNL houses Titan, a Cray XK7 supercomputer, which achieved the number one spot on the TOP500 at 17.59 petaflop/s in November 2012, making it the fastest supercomputer in the world. The National Institute for Computational Sciences (NICS), a joint UT/ORNL computing facility in Oak Ridge, houses Kraken, UT's Cray XT5 system which is one of the world's fastest open-science supercomputers. NICS is also home to Beacon, an Appro Xtreme-X Supercomputer which topped the Green500 list in November 2012, making it the most energy efficient supercomputer in the world.

GRID RESOURCES

With the continuing trend of high performance grid and cloud computing, it is important for ICL to have access to these types of infrastructures in order to test and implement our software packages on grid hardware and virtualized environments. In keeping with this goal, research staff members at ICL have access to grid resources all over the US, as well as in Europe, including XSEDE, Grid5000, and FutureGrid.

NVIDIA
CUDA™ RESEARCH CENTER

In 2009, the Innovative Computing Laboratory was designated a CUDA Center of Excellence (CCOE) by NVIDIA Corporation, a world-wide leader in technologies for visual computing and inventor of the graphics processing unit (GPU). This award led to the establishment of a productive long-term collaboration between ICL and NVIDIA. As part of the collaboration and CCOE designation, ICL has continuously received state-of-the-art hardware, financial support, and other resources from NVIDIA, while our accomplishments in the development of numerical linear algebra libraries for CUDA-based architectures have made a major impact on the nation's software cyberinfrastructure, and thus have further established the use of NVIDIA technologies in the forefront of parallel computing.

Joining a very small and select group of CUDA CCOEs such as labs at Harvard University, the University of Utah, and the University of Illinois at Urbana-Champaign, UTK's CCOE focuses on the development of numerical linear algebra libraries for CUDA-based hybrid architectures. Our work on the Matrix Algebra on GPU and Multicore Architectures (MAGMA) project, as well as on the Keeneland project, further enables and expands our CUDA-based numerical libraries development efforts, especially in the general area of high-performance scientific computing.

LOCAL LEADERSHIP

Over the past 15 years, UTK, ORNL, and the East Tennessee region have become renowned as leaders in high performance scientific computing. This achievement is due in no small part by the efforts of ICL, both through the consistent high quality of its research and through its collaboration, education, and outreach efforts. When the University established the Center for Information Technology Research (CITR) in 2001, with Professor Jack Dongarra as its director, its mission was to help provide technical and administrative support for research and outreach efforts on campus, including ICL.

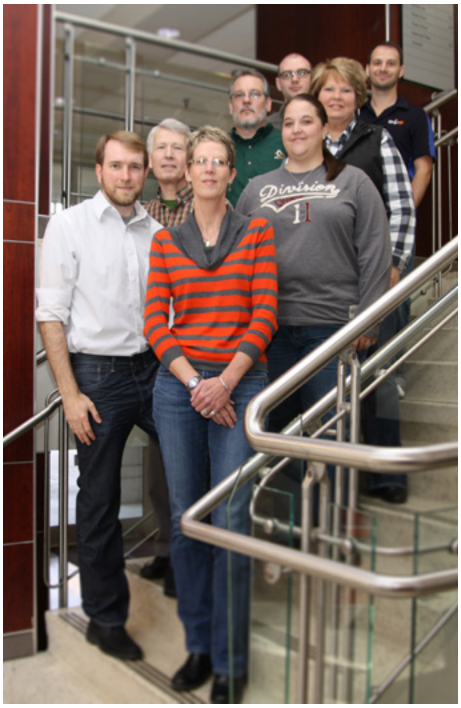

CITR Staff
(First Row) David Rogers, Tracy Rafferty;
(Second Row) Terry Moore, Leighanne Sisk;
(Third Row) Don Fike, Teresa Finchum;
(Fourth Row) Sam Crawford, Paul Peltz

CITR
CENTER FOR INFORMATION TECHNOLOGY RESEARCH

The Center for Information Technology Research was established in 2001 to drive the growth and development of leading edge information technology research at the University of Tennessee. CITR's first objective is to build up a thriving, well-funded community in basic and applied information technology research at UT in order to help the university capitalize on the rich supply of research opportunities that now exist in this area. As part of this goal, CITR staff members currently provide primary administrative and technical support for ICL, helping maintain the lab's status as a world leader in high performance and scientific computing research. CITR has also provided secondary support for other UT research centers.

CITR's second objective is to grow an interdisciplinary Computational Science program as part of the university curriculum. To this end, CITR helped establish the Interdisciplinary Graduate Minor in Computational Science to offer UT graduate students an opportunity to acquire the balanced package of knowledge and skills required for today's computationally intensive research methods. CITR is also the sole provider of administrative support for the IGMCS.

FIND OUT MORE AT http://citr.eecs.utk.edu/

To develop the kind of human capital it takes to keep our region at the forefront of HPC, and maintain UTK as a world leader in computational science and research, in 2007 CITR spearheaded the effort to create, and now provides primary support for, the Interdisciplinary Graduate Minor in Computational Science (IGMCS), UTK's only campus wide program in Computational Science. We expect ICL and CITR to continue to play a central role in the growth and evolution of Computational Science research and education at UTK for the foreseeable future.

IGMCS

INTERDISCIPLINARY GRADUATE MINOR IN COMPUTATIONAL SCIENCE

Addressing the need for a new educational strategy in Computational Science, the Center for Information Technology Research worked with faculty and administrators from several departments and colleges in 2007 to help establish a new university-wide program that supports advanced degree concentrations in this critical new field across the curricula. Under the Interdisciplinary Graduate Minor in Computational Science, students pursuing advanced degrees in a variety of fields of science and engineering are able to extend their education with special courses of study that teach them both the fundamentals and the latest ideas and techniques from this new era of information intensive research.

Computational Science integrates elements that are normally studied in different parts of the traditional curriculum, but which are not fully covered or combined by any one of them. As computational power continues to increase and data storage costs decrease, the potential for new discoveries using Computational Science is greater than ever. And as more academic disciplines begin to realize and exploit the incredible benefits Computational Science provides, the IGMCS program is expected to grow by adding new disciplines, new courses, and new faculty. As of late 2012, there were 17 departments from four UT colleges contributing more than 100 courses to the program.

FIND OUT MORE AT http://igmcs.utk.edu/

PEOPLE

As the landscape in high performance computing continues to rapidly evolve, remaining at the forefront of discovery requires great vision and skill. To address this evolution and to remain a leader in innovation, we have assembled a staff of top researchers from all around the world, who apply a variety of novel and unique approaches to the challenges and problems inherent in world-class, scientific computing.

As part of an engineering college at a top 50 public research university, we have a responsibility to combine exemplary teaching with cutting-edge research. As such, we regularly employ more than a dozen bright and motivated graduate and undergraduate students. During the fall of 2011, we recruited 8 graduate students from institutions all over the world, including the University of Tennessee's own Electrical Engineering and Computer Science department. We have been, and will continue to be, very proactive in securing internships and assistantships for students who are hardworking and willing to learn.

CURRENT STAFF AND STUDENTS

Wesley Bland
Graduate Research Assistant

George Bosilca
Research Assistant Professor

Aurelien Bouteiller
Research Scientist II

Chongxiao Cao
Graduate Research Assistant

Sam Crawford
Information Specialist I

Yuanshun Dai
Assistant Professor

Anthony Danalis
Research Scientist II

Simplice Donfack
Post Doc. Research Associate

Tingxing Dong
Graduate Research Assistant

Jack Dongarra
Univ. Distinguished Professor

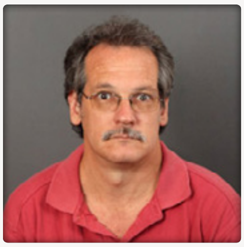

Don Fike
IT Specialist III

Teresa Finchum
Administrative Specialist II

Mark Gates
Research Scientist I

Peter Gaultney
Graduate Research Assistant

Azzam Haidar
Research Scientist I

Blake Haugen
Graduate Research Assistant

Thomas Herault
Research Scientist II

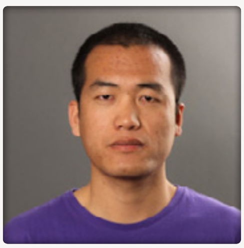

Yulu Jia
Graduate Research Assistant

Matt Johnson
Graduate Research Assistant

Vijay Joshi
Graduate Research Assistant

CURRENT STAFF AND STUDENTS
CONTINUED

Khairul Kabir
Graduate Research Assistant

Jakub Kurzak
Research Director

Julie Langou
Research Leader

Piotr Luszczek
Research Director

Teng Ma
Graduate Research Assistant

Heike McCraw
Research Scientist II

Terry Moore
Associate Director

Stephanie Moreaud
Post Doc. Research Associate

John Nelson
Graduate Research Assistant

Paul Peltz
IT Administrator III

Tracy Rafferty
Program Manager

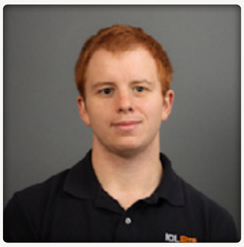

James Ralph
Research Associate II

David Rogers
IT Specialist III

Leighanne Sisk
Administrative Specialist I

Dan Terpstra
Research Leader II

Stanimire Tomov
Research Director

Wei Wu
Graduate Research Assistant

Ichitaro Yamazaki
Research Scientist I

Asim YarKhan
Senior Research Associate

VISITORS

Since ICL was founded, we have routinely hosted many visitors, some who stay briefly to give seminars or presentations, and others who remain with us for as long as a year collaborating, teaching, and learning. By collaborating with researchers from around the globe, we are able to leverage an immense array of intellectual resources. For this reason, our list of research collaborators and partners continues to grow. These relationships present enormous opportunities to host and work with top minds within the global HPC community.

2012 VISITORS

Emmanuel Agullo
INRIA
France

Ahmad Ahmad
King Abdullah University of Science and Technology
Saudi Arabia

Dorian Arnold
Univeristy of New Mexico
USA

Guillaume Aupy
ENS Lyon
France

Marc Baboulin
Université Paris-Sud
France

Giuseppe Bruno
Bank of Italy
Italy

Cedric Castagnede
INRIA Bordeaux
France

Vincent Cohen-Addad
ENS Lyon
France

Dave Cronk
Lockheed-Martin
USA

Wilfried Gansterer
University of Vienna
Austria

Laura Grigori
INRIA
France

Lou Gross
University of Tennessee
NIMBioS
USA

Hidehiko Hasegawa
University of Tsukuba
Japan

Julien Herrmann
ENS Lyon
France

Atsushi Hori
RIKEN
Japan

Mitch Horton
Georgia Tech
USA

Yutaka Ishikawa
University of Tokyo
Japan

Branislav Jansik
Aarhus University
Denmark

Xavier Lacoste
INRIA Bordeaux
France

Mabule Mabakane
University of Cape Town
South Africa

Gabriel Marin
Oak Ridge National Laboratory
USA

Guillaume Mercier
University of Bordeaux
France

Cleve Moler
MathWorks
USA

Daichi Mukunoki
University of Tsukuba
Japan

Nicholas Nagle
University of Tennessee
Department of Geography
USA

James Plank
University of Tennessee
EECS
USA

Pierre Ramet
INRIA
France

Yves Robert
ENS Lyon
France

Mitsuhisa Sato
University of Tsukuba
Japan

Volodymyr Turchenko
Ternopil National Economic University
Ukraine

Rick Weber
University of Tennessee
EECS
USA

Erlin Yao
Institute of Computing Technology, Chinese Academy of Sciences
China

Omar Zenati
ENSEIRB-MATMECA
France

ICL ALUMNI

Carolyn Aebischer 1990-1993
Bivek Agrawal 2004-2006
Sudesh Agrawal 2004-2006
Emmanuel Agullo 2009
Ahmad Abdelfattah Ahmad 2012
Jennifer Allgeyer 1993
Wes Alvaro 2007-2011
Ed Anderson 1989-1991
Daniel Andrzejewski 2007
Thara Angskun 2003-2007
Hartwig Anzt 2011
Papa Arkhurst 2003
Dorian Arnold 1996-2001
Cedric Augonnet 2010
Marc Baboulin 2008
Zhaojun Bai 1990-1992
Ashwin Balakrishnan 2001-2002
Richard Barrett 1992-1994
Alex Bassi 2000-2001
David Battle 1990-1992
Micah Beck 2000-2001
Daniel Becker 2007
Dulceneia Becker 2010-2012
Adam Beguelin 1991
Annamaria Benzoni 1991
Tom Berry 1991
Vincent Berthoux 2010
Scott Betts 1997-1998
Nikhil Bhatia 2003-2005
Noel Black 2002-2003
Laura Black 1996
Susan Blackford 1989-2001
Kartheek Bodanki 2009
David Bolt 1991
Fernando Bond 1999-2000
Carolyn Bowers 1992
Barry Britt 2007-2009
Randy Brown 1997-1999
Cynthia Browne 2005
Murray Browne 1998-1999
Bonnie Browne 2011-2012
Antonin Bukovsky 1998-2003
Greg Bunch 1995
Alfredo Buttari 2004-2007
Giuseppe Bruno 2001
Anthony Canino 2012
Domingo Gimenez Canovas 2001
Henri Casanova 1995-1998
Cedric Castagnede 2012
Ramkrishna Chakrabarty 2005
Sharon Chambers 1998-2000
Zizhong Chen 2001-2006
Jaeyoung Choi 1994-1995
Wahid Chrabakh 1999
Eric Clarkson 1998
Andy Cleary 1995-1997
Michelle Clinard 1989-1991
Vincent Cohen-Addad 2012
Matthias Colin 2004
Charles Collins 2012
Tom Cortese 2002-2009
Camille Coti 2007
Jason Cox 1993-1997
David Cronk 1999 - 2010
Javier Cuenca 2003
Manoel Cunha 2006
Cricket Deane 1998-1999
Remi Delmas 2006
Frederic Desprez 1994-1995
Jin Ding 2003
Jun Ding 2001-2003
Ying Ding 2000-2001
Martin Do 1993-1994
Leon Dong 2000-2001
Nick Dongarra 2000
David Doolin 1997
Andrew Downey 1998-2003
Mary Drake 1989-1992
Julio Driggs 2002-2004
Brian Drum 2001-2004
Peng Du 2005-2012
Eduardo Echavarria 2005
Victor Eijkhout 1992-2005
Brett Ellis 1995-2005
Shawn Ericson 2004
Zachary Eyler-Walker 1997-1998
Lisa Ezzell 2003-2004
Christoph Fabianek 2003
Graham Fagg 1996-2006
Mathieu Faverge 2010-2012
Shengzhog Feng 2005-2006
Salvatore Filippone 2004
Anna Finchum 2010
Mike Finger 1997
Markus Fischer 1997-1998
Len Freeman 2009
Xiaoquan Fu 2003-2004
Erika Fuentes 2003-2007
Karl Fuerlinger 2006-2008
Megan Fuller 2006
Edgar Gabriel 2003-2004
Tracy Gangwer 1992-1993
Lynn Gangwer 2000-2001
Nathan Garner 2001-2006
Kelley Garner 1998
Tina Garrison 1991
Adriana Garties 2011
Christoph Geile 2008
Jean Patrick Gelas 2001
Boris Gelfend 1993
Jonathan Gettler 1996
Eric Greaser 1993
Stan Green 1992-1996
Alice Gregory 2004-2006
Jason Gurley 1997-1998
Bilel Hadri 2008-2009
Hunter Hagewood 2000-2001
Christian Halloy 1996-1997
Sven Hammarling 1996-1997
J. Mike Hammond 1994-1995
Satomi Hasegawa 1995-1996
Hidehiko Hasegawa 1995-1996
Chris Hastings 1996
David Henderson 1999-2001
Greg Henry 1996
Julien Herrmann 2011-2012
Holly Hicks 1993-1994
Alexandra Hicks-Hardiman 2009
Sid Hill 1996-1998
Tomoyuki Hiroyasu 2002-2003
George Ho 1998-2000
Josh Hoffman 2008-2010
Jeff Horner 1995-1999
Mitch Horton 2010-2012
Yan Huang 2000-2001
Aurelie Hurault 2009
Chris Hurt 2002
Paul Jacobs 1992-1995
Emmanuel Jeannot 2001-2006
Weizhong Ji 1999-2000
Weicheng Jiang 1992-1995
Song Jin 1997-1998
Patrick Johansson 2001
Aral Johnson 2009
Sean Jolly 1997-1998
Jan Jones 1992-2008
Kim Jones 1996-1997
Venkata Kakani 2007
Ajay Kalhan 1995
Balajee Kannan 2001
Madhuri Kasam 2007-2008
Kiran Kasichayanula 2010-2012

Since its inception, ICL has attracted many students, post-doctoral researchers, and professors from a variety of backgrounds and academic disciplines. Many of these experts came to UT specifically to work with Dr. Dongarra, beginning a long list of top research talent to pass through ICL and move on to make exciting contributions at other institutions and organizations.

Ajay Katta 2010
David Katz 2002
Joshua Kelly 2000-2001
Supriya Kilambi 2008
Myung Ho Kim 2005-2006
Youngbae Kim 1992-1996
Jenya Kirshtein 2008
Michael Kolatis 1993-1996
Chandra Krintz 1999-2001
Tilman Kuestner 2010
Krerkchai Kusolchu 2010
Coire Kyle 2005
Amanda Laake 2003-2004
Xavier Lacoste 2012
Julien Langou 2003-2006
Jeff Larkin 2003-2005
Brian LaRose 1990-1992
Frank Lauer 2010
DongWoo Lee 2000-2002
Tracy Lee 1996-2012
Klaudia Leja 2008
Pierre Lemarinier 2008-2010
Todd Letsche 1993-1994
Sharon Lewis 1992-1995
Xiang Li 2001
Yinan Li 2006-2008
Weiran Li 2002
Chaoyang Liu 2000
Kevin London 1996-2005
Matt Longley 1999
Hatem Ltaief 2008-2011
Daniel Lucio 2008
Richard Luczak 2000-2001
Robert Manchek 1990-1996
Tushti Marwah 2004
Donald McCasland 1994
Paul McMahan 1994-2000
Eric Meek 2003-2006
James Meyering 1991-1992

Jeremy Millar 1998-2002
Michelle Miller 1999-2003
Cindy Mitchell 2001-2002
Stuart Monty 1993
Eric Moore 2000
Keith Moore 1987-2007
Shirley Moore 1993-2012
Robert Morgan 1990-1991
Kishan Motheramgari 1997
Steven Moulton 1991-1993
Phil Mucci 1994-2004
Daichi Mukunoki 2012
Matthew Nabity 2008
Shankar Narasimhaswami 2004-2005
Rajib Nath 2008-2010
Fernando Navarro 2009
Donnie Newell 2010
Peter Newton 1994-1995
Jonas Nilsson 2001
Jakob Oestergaard 2000
Caroline Papadopoulos 1997-1998
Leelinda Parker 2002
Dilip Patlolla 2007-2008
Andy Pearson 1989-1991
Theresa Pepin 1994
Antoine Petitet 1993-2001
Peter Pham 2012
Vlado Pjesivac 2008
Jelena Pjesivac-Grbovic 2003-2007
James S. Plank 1991-1992
Ciara Proctor 2008
Tim Poore 2009
Roldan Pozo 1992-1994
Farzona Pulatova 2005-2006
Martin Quinson 2001
Tammy Race 1999-2001
Ganapathy Raman 1998-2000
Kamesh Ramani 2003
Mei Ran 1999-2004

Arun Rattan 1997
Sheri Reagan 1995-1996
Mike Reynolds 1994
George Rhinehart 2012
Jon Richardson 1990-1991
Ken Roche 1999-2004
Andrew Rogers 1997-1999
Tom Rothrock 1997-1998
Tom Rowan 1993-1997
Narapat (Ohm) Saengpatsa 2011
Kiran Sagi 2001-2005
Evelyn Sams 1998-1999
Ken Schwartz 1992-1993
Keith Seymour 1994-2009
Farial Shahnaz 2001
Brian Sheely 2009-2010
Zhiao Shi 2001-2007
Sergei Shinkarev 2005-2007
Majed Sidani 1991-1992
Shilpa Singhal 1996-1998
Matt Skinner 2008
Jonte Smith 2008
Peter Soendergaard 2000
Raffaele Solca 2012
Gwang Son 2007-2009
Fengguang Song 2003-2012
Thomas Spencer 1999-2001
Erich Strohmaier 1995-2001
Xiaobai Sun 1995
Martin Swany 1996-1999
Daisuke Takahashi 2002
Judi Talley 1993-1999
Ronald Tam 2009
Yuan Tang 2005-2006
Yusuke Tanimura 2003
Keita Teranishi 1998
Joe Thomas 2002-2009
John Thurman 1998-1999
Francoise Tisseur 1997

Jude Toth 1993-1994
Bernard Tourancheau 1993-1994
Malcolm Truss 2008
Lauren Vaca 2004
Sathish Vadhiyar 1999-2003
Robert van de Geijn 1990-1991
Chad Vawter 1995
Eugene Vecharynski 2008
Scott Venckus 1993-1995
Antoine Vernois 2004
Reed Wade 1990-1996
Michael Walters 2001-2005
Mike Waltz 1999
Robert Waltz 1990-1991
Jerzy Wasniewski 2000
Vince Weaver 2010-2012
Scott Wells 1997-2010
David West 1990-1992
R. Clint Whaley 1991-2001
Jody Whisnant 1997-1998
James White 1999
Scotti Whitmire 1995-1996
Susan Wo 2000-2001
Felix Wolf 2003-2005
Jiayi Wu 2004-2007
Qiu Xia 2004-2005
Tinghua Xu 1998-2000
Tao Yang 1999
Jin Yi 2009-2010
Haihang You 2004-2009
Lamia Youseff 2007
Brian Zachary 2009-2010
Omar Zenati 2012
Yuanlei Zhang 2001-2005
Junlong Zhao 2002
Yong Zheng 2001
Luke Zhou 2000-2001
Min Zhou 2002-2004

SPONSORS

For more than 23 years, our knowledge and hard work have earned the trust and support of many agencies and organizations that have funded, and continue to fund, our efforts. Without them we simply would not be able to conduct cutting-edge research. The main source of support has been federal agencies that are charged with allocating public research funding. Therefore, we acknowledge the following agencies for supporting our efforts, both past and present:

In addition to the support of the federal government, we have solicited strong support from private industry, which has also played a significant role in our success and growth. Some organizations have targeted specific ICL projects, while others have made contributions to our work that are more general and open-ended. We gratefully acknowledge the following vendors for their generosity and support:

Administrative and technical support for ICL is provided by the Center for Information Technology Research (CITR).

Financial support for CITR is provided by the University of Tennessee.

INNOVATIVE
COMPUTING LABORATORY
THE UNIVERSITY of TENNESSEE
Department of Electrical Engineering and Computer Science

Suite 203 Claxton
1122 Volunteer Blvd
Knoxville, TN 37996-3450

FIND OUT MORE AT

http://icl.utk.edu/

INNOVATIVE
COMPUTING LABORATORY
THE UNIVERSITY of TENNESSEE

CITR CENTER FOR INFORMATION TECHNOLOGY RESEARCH